Dynamic Role Authorization in Multiparty Conversations

Silvia Ghilezan
Univerzitet u Novom Sadu, Serbia

Svetlana Jakšić
Univerzitet u Novom Sadu, Serbia

Jovanka Pantović
Univerzitet u Novom Sadu, Serbia

Jorge A. Pérez
University of Groningen, The Netherlands

Hugo Torres Vieira
LaSIGE, Faculdade de Ciências, Universidade de Lisboa, Portugal

Protocol specifications often identify the *roles* involved in communications. In multiparty protocols that involve task delegation it is often useful to consider settings in which different sites may act on behalf of a single role. It is then crucial to control the roles that the different parties are authorized to represent, including the case in which role authorizations are determined only at runtime. Building on previous work on conversation types with flexible role assignment, here we report initial results on a typed framework for the analysis of multiparty communications with dynamic role authorization and delegation. In the underlying process model, communication prefixes are annotated with role authorizations and authorizations can be passed around. We extend the conversation type system so as to statically distinguish processes that never incur in authorization errors. The proposed static discipline guarantees that processes are always authorized to communicate on behalf of an intended role, also covering the case in which authorizations are dynamically passed around in messages.

1 Introduction

Different concepts of role-based performance can be found in modern distributed information systems, ranging from access control to structured interactions in communication-centred systems. These concepts are typically grounded on the assumption that distinct participants (e.g., users at different physical locations) may belong to the same role, and that a single participant may belong to (or implement) several different roles. Each role is associated with a set of permissions (e.g., privileges to access data or perform some action), thus enforcing an assignment of permissions to involved participants. In the case of multi-party interactions, a participant can use a role for communication only if the role is authorized for the particular action. As an example, consider the scenario of an electronic submission system, in which (confidential) paper submissions should be available only to only authorized participants. In this scenario, editors typically rely on other participants who may act as reviewers. As such, any participant should be aware of the possibility of being appointed as reviewer. Also, a participant should be able to act as a reviewer only when she is officially authorized by the editor—this means, in particular, that unauthorized participants must not be able to read a submission. Furthermore, the exchanges determining an authorization should be part of the predefined protocols between the editor and the reviewer-to-be.

In this paper, we consider the issue of enhancing multiparty communications with *dynamic role authorizations*. Our starting point is the typed framework given in [1], based on *conversation types*, in which roles are flexibly assigned to participants that can act on their behalf. While expressive, the model in [1] does not check whether a given assigned role is indeed authorized to perform a particular action. The model that we propose here addresses this shortcoming, explicitly tracking the presence of (un)authorized actions. Our typed model also enables the exchange of authorizations along communication actions. Hence, participants may dynamically obtain authorizations to act on behalf of a role. We view our contribution as a first step in modeling and analyzing dynamic role-based communication and

Marco Carbone (Ed.): Third Workshop on
Behavioural Types (BEAT 2014)
EPTCS 162, 2014, pp. 1–8, doi:10.4204/EPTCS.162.1

© Ghilezan, Jakšić, Pantović, Pérez, Vieira
This work is licensed under the
Creative Commons Attribution License.

authorization, focusing on the identification of the basic ingredients that should be added on top of an existing framework in order to address the problem.

We present the main highlights of our model by formalizing the submission system. Let us assume the set of roles $\{\texttt{professor}, \texttt{student}, \texttt{reviewer}, \texttt{editor}\}$ and the following global specification:

$$\begin{array}{l}(\texttt{editor} \to \texttt{professor}) : \mathit{auth}1(\texttt{reviewer}).\\ \quad(\texttt{professor} \to \texttt{student}) : \mathit{auth}2(\texttt{reviewer}).\\ \quad\quad(\texttt{editor} \to \texttt{reviewer}) : \mathit{extend}().\\ \quad\quad\quad(\texttt{student} \to \texttt{professor}) : \mathit{report}().\\ \quad\quad\quad\quad(\texttt{reviewer} \to \texttt{editor}) : \mathit{final}().\texttt{end}.\end{array}$$

Above, role `editor` is allowed to send authorizations for the role `reviewer`. The global specification says that the `editor` authorizes the `professor` to act as `reviewer`, which is followed by passing the authorization for the role `reviewer` from the `professor` to the `student`. The `reviewer` gets a deadline extension from the `editor`, then `student` sends the report to the `professor`. Finally, the `reviewer` sends a final decision to the `editor`. We may implement this specification as the process

$$\mathit{Sys} = (\nu\,\texttt{subm})\texttt{journal}_{\lceil\texttt{editor}\rceil}!\mathit{paper}(\texttt{subm}).P' \mid \texttt{journal}_{\lceil\texttt{professor}\rceil}?\mathit{paper}(a).Q' \mid R$$

where `subm` denotes a channel and processes P', Q', and R are defined as:
$P' = \texttt{subm}_{\lceil\texttt{editor}\rceil}!\mathit{auth}1(\lceil\texttt{reviewer}\rceil).P''$
$P'' = \texttt{subm}_{\lceil\texttt{editor}\rceil}!\mathit{extend}().\texttt{subm}_{\lceil\texttt{editor}\rceil}?\mathit{final}().0$
$Q' = a_{\lceil\texttt{professor}\rceil}?\mathit{auth}1(\texttt{reviewer}).Q''$
$Q'' = \texttt{assist}_{\lceil\texttt{professor}\rceil}!\mathit{read}(a).a_{\lceil\texttt{professor}\rceil}!\mathit{auth}2(\lfloor\texttt{reviewer}\rfloor).$
$\quad a_{\lceil\texttt{professor}\rceil}?\mathit{report}().a_{\lfloor\texttt{reviewer}\rfloor}!\mathit{final}().0$
$R = \texttt{assist}_{\lceil\texttt{student}\rceil}?\mathit{read}(b).b_{\lceil\texttt{student}\rceil}?\mathit{auth}2(\texttt{reviewer}).b_{\lfloor\texttt{reviewer}\rfloor}?\mathit{extend}().b_{\lceil\texttt{student}\rceil}!\mathit{report}().0$

In our process model, each communication prefix α is decorated with either $\lceil r\rceil$ (i.e., role r is authorized to perform α) or $\lfloor r\rfloor$ (i.e., role r is *not* authorized to perform α). These decorations define fine-grained specifications of (un)authorized communication actions. The three subprocesses of *Sys* formalize the behavior of the editor, the professor, and the student, respectively. The first subprocess creates a fresh channel `subm` which is passed (on behalf of `editor`) to the second subprocess (that receives it as `professor`) over the message *paper* on channel `journal`. Process *Sys* then reduces to

$$(\nu\,\texttt{subm})(\texttt{subm}_{\lceil\texttt{editor}\rceil}!\mathit{auth}1(\lceil\texttt{reviewer}\rceil).P'' \mid \texttt{subm}_{\lceil\texttt{professor}\rceil}?\mathit{auth}1(\texttt{reviewer}).Q''\{\texttt{subm}/a\} \mid R)$$

Here the process authorized as `editor` sends authorization for the role `reviewer` to the process acting on behalf of `professor`. After interaction, the role `reviewer` will become authorized in $Q''\{\texttt{subm}/a\}$, so the second subprocess will reduce to

$$\begin{array}{l}\texttt{assist}_{\lceil\texttt{professor}\rceil}!\mathit{read}(\texttt{subm}).\texttt{subm}_{\lceil\texttt{professor}\rceil}!\mathit{auth}2(\lceil\texttt{reviewer}\rceil).\\ \quad\quad\texttt{subm}_{\lceil\texttt{professor}\rceil}?\mathit{report}().\texttt{subm}_{\lceil\texttt{reviewer}\rceil}!\mathit{final}().0\end{array}$$

Continuing along these lines, process R joins the conversation on channel `subm`, gets authorization for `reviewer`, receives *extend* as `reviewer`, sends *report* as `student`, and finally the second subprocess sends *final* decision, as `reviewer`, to the process acting as `editor`. It is worth observing that the initial specification for the student (cf. process R) is authorized to act as `student` but it is not authorized to act as `reviewer`. The required authorization to access and review the submission should result as a

consequence of an interaction with the process realizing the behavior of the professor. That is, previous communication actions directly determine current authorization privileges for interacting partners. As such, the issue of ensuring consistent conversations is tightly related to issues of role authorization and deauthorization. To address this combination of issues, the type discipline that we present here ensures that structured multiparty conversations are consistent with respect to both global protocols and requirements of dynamic role authorization (cf. Corollary 1).

This paper is organized as follows. In § 2, we define our process language and illustrate further the intended model of dynamic role authorization. § 3 presents the type system and the properties for well-typed processes. Finally, in § 4, we comment on related works and discuss open problems.

2 Process Language

Syntax. We consider a synchronous π-calculus [10] extended with labelled communications and prefixes for authorization sending and receiving. Let $\mathscr{L}$, $\mathscr{R}$, and $\mathscr{N}$ be infinite base sets of *labels*, *roles*, and *channels*, respectively. We use $l,\ldots$ to range over $\mathscr{L}$; $r,s,\ldots$ to range over $\mathscr{R}$; and $a,b,c,\ldots$ to range over $\mathscr{N}$. A role r can be qualified as *authorized* (denoted $\lceil r \rceil$) or as *unauthorized* (denoted $\lfloor r \rfloor$). We write $\lfloor r \rceil$ to denote a role r with some qualification, and use ρ,σ to range over $\lfloor r \rceil$ for some unspecified r.

As motivated above, each communication prefix in our calculus is decorated by a qualified role for performing the associated action. Intuitively, prior to execution not all roles have to be authorized; we expect unauthorized roles may have the potential of becoming authorized as the structured interactions take place. In fact, we expect all actions of the system to be associated to authorized roles; top-level prefixes on unauthorized roles are regarded as *errors*. To enable dynamic role authorization, our model allows for the exchange of the authorization on a role. Formally, the syntax of processes is given by:

$$\begin{array}{rcl} P,Q & ::= & 0 \;\mid\; P \mid Q \;\mid\; (\nu a)P \;\mid\; \alpha.P \\ \alpha & ::= & a_\rho!l(b) \;\mid\; a_\rho?l(b) \;\mid\; a_\rho!l(\sigma) \;\mid\; a_\rho?l(r) \qquad\qquad \rho ::= \lceil r \rceil \mid \lfloor r \rfloor \end{array}$$

Constructs for inaction (0), parallel composition ($P \mid Q$), and restriction ($(\nu a)P$) are standard. We write $(\nu a_1,\ldots,a_n)$ as a shorthand for $(\nu a_1)\cdots(\nu a_n)$. Also, we write $\widetilde{a}$ to stand for the sequence of names $a_1,\ldots,a_n$. To define communication of channels and authorizations, our language has four kinds of prefixes, denoted α. Each prefix is associated to a ρ. Intuitively, a prefix associated to $\lceil r \rceil$, is said to be *authorized* to perform the associated action under role r. A prefix associated to $\lfloor r \rfloor$ is not authorized to perform the corresponding action as r; but it may be the case that such prefixes are dynamically authorized via communication. The intuitive semantics for prefixes follows:

- $a_\rho!l(b)$ expresses sending of name b, in labelled message l, along channel a, under qualified role ρ;
- $a_\rho?l(b)$ expresses receiving of name b, in labelled message l, along channel a, under qualified role ρ.

These two prefixes are taken from [1], here extended in with authorization control via role qualification. The second pair of prefixes is new to our calculus:

- $a_\rho!l(\sigma)$ expresses sending of the *qualified role* σ, in labelled message l, along channel a, under qualified role ρ;
- $a_\rho?l(r)$ expresses receiving authorization for role r, in labelled message l, along channel a, under qualified role ρ.

[R-COMM] $a_{\lceil s\rceil}!l(b).P \mid a_{\lceil r\rceil}?l(c).Q \to P \mid Q\{b/c\}$

[R-AUTH] $a_{\lceil s\rceil}!l(\lceil q\rceil).P \mid a_{\lceil r\rceil}?l(q).Q \to P \mid Q\{a : \lceil q\rceil/\lfloor q\rfloor\}$

[R-PAR] $P \to P' \;\Rightarrow\; P \mid Q \to P' \mid Q$

[R-RESTRICTION] $P \to Q \;\Rightarrow\; (\nu a)P \to (\nu a)Q$

[R-STRUCT] $P \equiv P', P' \to Q', Q' \equiv Q \;\Rightarrow\; P \to Q$

Figure 1: Reduction relation

Operational Semantics. The process semantics is defined via a reduction relation, which is defined in Figure 1 and denoted $\to$. Reduction is closed under static contexts and *structural congruence*, denoted $\equiv$ and defined in standard lines (cf. [10]). To support communication of qualified roles, we use a form of substitution denoted by $a : \{\sigma/\overline{\sigma}\}$, representing the substitution $\{\sigma/\overline{\sigma}\}$ applied only on channel a. In turn, this includes two substitutions, for prefix subjects and qualified roles occurring as communication objects, respectively: $\{a_\sigma/a_{\overline{\sigma}}\}$ and $\{a_\rho!l(\sigma)/a_\rho!l(\overline{\sigma})\}$. In Figure 1 rules [R-Comm] and [R-Auth] specify synchronizations: two processes can exchange a message (l) on a channel (a) only under authorized roles (denoted $\lceil s\rceil, \lceil r\rceil$ in the rules). Using rule [R-Auth], a process can authorize another process to act under a role (q) only if the first process has permission to do such an authorization ($\lceil q\rceil$) and the second process is asking for authorization of the same role (q). We use $\to^*$ to denote the reflexive and transitive closure of $\to$. The following definitions are key to syntactically distinguish authorization errors:

Definition 1 (Unauthorized Prefix / Errors) *Let α and P be a prefix and a process as defined above.*

- *We say α is* unauthorized *if its subject is associated to an unauthorized role or its output object is an unauthorized role, i.e., if $\alpha = a_{\lfloor r\rfloor}?l(b)$, $\alpha = a_{\lfloor r\rfloor}!l(b)$, $\alpha = a_{\lfloor r\rfloor}?l(s)$, $\alpha = a_{\lfloor r\rfloor}!l(\lceil s\rceil)$, or $a_\sigma!l(\lfloor s\rfloor)$, for some a, r, l, b, s, σ. We write $\lfloor\alpha\rfloor.P$ instead of $\alpha.P$ whenever α is unauthorized.*
- *We say that P is an* authorization error *if $P \equiv (\nu\widetilde{a})(\lfloor\alpha\rfloor.Q \mid R)$, for some $\widetilde{a}, \alpha, Q, R$.*

Notice that an authorization error is a "stuck process" according to our semantics, i.e., a process which cannot synchronize since it does not have the required authorization to do so.

Example 1 *To illustrate reduction and authorization errors, consider processes P and Q below:*

$$\begin{aligned} P &= a_{\lceil r\rceil}?l_2(c).c_{\lceil r\rceil}?l_1(s).c_{\lfloor s\rfloor}?l_3().0 \mid (\nu b)a_{\lceil q\rceil}!l_2(b).b_{\lceil q\rceil}!l_1(\lceil s\rceil).b_{\lceil q\rceil}!l_3().0 \\ Q &= (\nu b)(b_{\lceil q\rceil}!l_1(\lceil s\rceil).a_{\lceil q\rceil}!l_2(b).b_{\lceil q\rceil}!l_3().0 \mid b_{\lceil r\rceil}?l_1(s).0) \mid a_{\lceil s\rceil}?l_2(c).c_{\lfloor s\rfloor}?l_3().0 \end{aligned}$$

In both P and Q, channel b is used according to the specification $(q \to r) : l_1(s).(q \to s) : l_3()$ which informally says first there is an interaction from role q to role r on message l_1, exchanging authorization on role s, followed by an interaction between q and s on l_3 (where, for the sake of simplicity, we omit contents of the message). We may infer the following reductions for P:

$$P \;\to\; (\nu b)(b_{\lceil r\rceil}?l_1(s).b_{\lfloor s\rfloor}?l_3().0 \mid b_{\lceil q\rceil}!l_1(\lceil s\rceil).b_{\lceil q\rceil}!l_3().0) \;\to\; (\nu b)(b_{\lceil s\rceil}?l_3().0 \mid b_{\lceil q\rceil}!l_3().0)$$

and so all actions are carried out on behalf of authorized roles. In contrast, we have that

$$Q \to^* (\nu b)(b_{\lceil q\rceil}!l_3().0 \mid b_{\lfloor s\rfloor}?l_3().0).$$

and so we infer that Q is ill-behaved since it reduces to an authorization error on role s.

As the previous example illustrates, there are processes which respect communication specifications but lead to authorization errors. The type system described in the following section addresses this issue.

$$\begin{array}{lcl} B & ::= & \mathtt{end} \mid B \mid B \mid \Diamond B \mid \mathtt{p}l(M).B \\ M & ::= & B \mid T \mid r \end{array} \qquad \begin{array}{lcl} T & ::= & l(B) \\ \mathtt{p} & ::= & !r \mid ?r \mid (r \to r) : \end{array}$$

Figure 2: Conversation types

3 Type System

We consider the conversation types language as presented in [1], extending message type M with the role r, so that we may capture role authorization passing. This is a rather natural extension, formally given by the syntax in Figure 2. Behavioral types B include: $\mathtt{end}$, which describes inaction; $B \mid B$, which allows to describe concurrent independent behavior; the sometime type $\Diamond B$, which says that behavior B may take place immediately or later on. Finally, a behavioral type $\mathtt{p}l(M).B$ describes a communication action identifying the role or roles involved, and whether the action is an input $?r$ or an output $!r$ or a message exchange $(r \to r)$:, a carried message type M and the behavior that is prescribed to take place after the communication action B.

We use behavioral types to specify the interactions in *linear* channels, where no communication races are allowed (which is to say that at any given moment, there can only be one matching pair of input/output actions). In our setting, where several parties may simultaneously use a channel, this *linear* communication pairing is ensured via message labels: at a given moment, there can be only one pair of processes able to exchange a labelled message. For shared channels, where communication races are allowed, we ensure consistent usage (but no structured protocol of interaction) via shared channel types T, which carry a (linear) behavioral type describing the usage delegated in the communication.

Message type M also captures the usages delegated in communications: in case M is a behavioral type B or a shared channel type T it describes how the receiving process uses the received channel; in case M identifies a role r then the message type captures an authorization delegation in the specified role.

Type environments separate linear and shared channel usages: Δ associates channels with (linear) behavioral types, given by $\Delta ::= \emptyset \mid \Delta, a : B$, whereas Γ associates channels with shared channel types, given by $\Gamma ::= \emptyset \mid \Gamma, a : T$. Typing rules rely on subtyping as well as on operators for apartness, well-formedness, and splitting of types. We refer to [1] for a details on these operations. While *type apartness* (#) refers to independent behaviors ensured via disjoint (message) label sets, *well-formedness* ensures that parallel behaviors are apart and that the sometime $\Diamond$ is not associated to message synchronizations — synchronizations are not allowed to take place sometime later, they are always specified to take place at a given stage in the protocol. The *subtyping relation* $<:$ allows for (some) behaviors which are prescribed to take place immediately to be used in contexts that expect such behaviors to take place sometime ($\Diamond$) further along. Type *splitting* supports the distribution of protocol "slices" among the participants in a conversation: we write $B = B_1 \circ B_2$ to say that behavior B may be split in behaviors B_1 and B_2 so that an overall behavior B may be distributed (e.g., in the two branches of a parallel composition). We remark that B_1 and B_2 may be further split so as to single out the individual contributions of each participant in a conversation, where decomposition is driven by the structure of the process in the typing rules.

We lift the split relation to Δ type environments in unsurprising lines: $\Delta, a : B = \Delta_1, a : B_1 \circ \Delta_2, a : B_2$ if $B = B_1 \circ B_2$ and $\Delta = \Delta_1 \circ \Delta_2$, and also $\Delta, a : B = \Delta_1, a : B \circ \Delta_2$ (and symmetrically) if $\Delta = \Delta_1 \circ \Delta_2$. In typing rules we write $\Delta_1 \circ \Delta_2$ to represent Δ (if there is such Δ) such that $\Delta = \Delta_1 \circ \Delta_2$.

A typing judgment is of the form $\Delta; \Gamma \vdash_\Sigma P$. The *authorization set* Σ is a subset of the direct product of the set of channel names and the set of roles, i.e., $\Sigma \subseteq \mathscr{N} \times \mathscr{R}$. The typing judgment states that the process P is well typed under Δ and Γ with roles from $\mathtt{pr}_2(\Sigma)$ (the projection on the second element of

$$\text{[T-end]}\quad \frac{}{\Delta_{\mathtt{end}};\Gamma\vdash_{\emptyset} 0} \qquad \text{[T-snew]}\quad \frac{\Delta;\Gamma,a:l(B)\vdash_{\Sigma} P}{\Delta;\Gamma\vdash_{\Sigma}(\nu a)P}$$

$$\text{[T-new]}\quad \frac{\Delta,a:B;\Gamma\vdash_{\Sigma} P \qquad \mathtt{matched}(B) \qquad a\notin \mathtt{pr}_1(\Sigma)}{\Delta;\Gamma\vdash_{\Sigma}(\nu a)P} \qquad \text{[TProc-par]}\quad \frac{\Delta_1;\Gamma\vdash_{\Sigma} P \qquad \Delta_2;\Gamma\vdash_{\Xi} Q}{\Delta_1\circ\Delta_2;\Gamma\vdash_{\Sigma\cup\Xi} P\mid Q}$$

$$\text{[Trole-in]}\quad \frac{\Delta\circ a:B;\Gamma\vdash_{\Sigma\cup\{(a,s)\}} P \quad ?r\,l(s).B<:B' \quad \Xi=\Sigma\cup\mathtt{unauth}(a,\lfloor r\rceil)}{\Delta\circ a:B';\Gamma\vdash_{\Xi} a_{\lfloor r\rceil}?l(s).P}$$

$$\text{[Trole-out]}\quad \frac{\Delta\circ a:B;\Gamma\vdash_{\Sigma} P \quad !r\,l(s).B<:B' \quad \Xi=\Sigma\cup\mathtt{unauth}(a,\lfloor r\rceil)\cup\mathtt{unauth}(a,\lfloor s\rceil)}{\Delta\circ a:B';\Gamma\vdash_{\Xi} a_{\lfloor r\rceil}!l(\lfloor s\rceil).P}$$

$$\text{[T-in]}\quad \frac{\Delta\circ a:B,b:B';\Gamma\vdash_{\Sigma} P \quad ?r\,l(B').B<:B'' \quad \Xi=\Sigma\cup\mathtt{unauth}(a,\lfloor r\rceil) \quad b\notin \mathtt{pr}_1(\Sigma)}{\Delta\circ a:B'';\Gamma\vdash_{\Xi} a_{\lfloor r\rceil}?l(b).P}$$

$$\text{[T-out]}\quad \frac{\Delta\circ a:B;\Gamma\vdash_{\Sigma} P \qquad !r\,l(B').B<:B'' \qquad \Xi=\Sigma\cup\mathtt{unauth}(a,\lfloor r\rceil)}{\Delta\circ a:B''\circ b:B';\Gamma\vdash_{\Xi} a_{\lfloor r\rceil}!l(b).P}$$

$$\text{[T-lsin]}\quad \frac{\Delta\circ a:B';\Gamma,b:T\vdash_{\Sigma} P \quad ?r\,l(T).B'<:B \quad \Xi=\Sigma\cup\mathtt{unauth}(a,\lfloor r\rceil) \quad b\notin \mathtt{pr}_1(\Sigma)}{\Delta\circ a:B;\Gamma\vdash_{\Xi} a_{\lfloor r\rceil}?l(b).P}$$

$$\text{[T-lsout]}\quad \frac{\Delta\circ a:B';\Gamma,b:T\vdash_{\Sigma} P \qquad !r\,l(T).B'<:B \qquad \Xi=\Sigma\cup\mathtt{unauth}(a,\lfloor r\rceil)}{\Delta\circ a:B;\Gamma,b:T\vdash_{\Xi} a_{\lfloor r\rceil}!l(b).P}$$

$$\text{[T-sin]}\quad \frac{\Delta,b:B;\Gamma,a:l(B)\vdash_{\Sigma} P \quad b\notin \mathtt{pr}_1(\Sigma)}{\Delta;\Gamma,a:l(B)\vdash_{\Sigma} a_{\lceil r\rceil}?l(b).P} \qquad \text{[T-sout]}\quad \frac{\Delta;\Gamma,a:l(B)\vdash_{\Sigma} P}{\Delta\circ b:B;\Gamma,a:l(B)\vdash_{\Sigma} a_{\lceil r\rceil}!l(b).P}$$

Figure 3: Typing rules. We define: $\mathtt{unauth}(a,\lfloor q\rfloor)=\{(a,q)\},\mathtt{unauth}(a,\lceil q\rceil)=\emptyset$

the pairs in Σ) appearing in P unauthorized on corresponding channels from $\mathtt{pr}_1(\Sigma)$.

Typing rules are presented in Figure 3. There are two rules that are specific to our model:

- [Trole-in] types authorization reception of the role s on the linear channel a, under the role $\lfloor r\rceil$, with the authorization set diminished by (a,s), and enlarged with (a,r) in case $\lfloor r\rceil$ is $\lfloor r\rfloor$.
- [Trole-out] types sending of the authorization $\lfloor s\rceil$ on linear channel a, under the role $\lfloor r\rceil$, with the authorization set enlarged with (a,r) or (a,s) in case $\lfloor r\rceil$ is $\lfloor r\rfloor$ or $\lfloor s\rceil$ is $\lfloor s\rfloor$.

Notice that in both rules the typing environment in the conclusion is split in a typing of a that specifies the reception of the authorization (up to subtyping).

All other rules are similar to the typing rules from [1], with the derivation of the novel decoration Σ as follows. Rule [T-end] states that a well-typed inactive process has no unauthorized roles and only `end` usages of linear channels (denoted by $\Delta_{\mathtt{end}}$). Rule [T-new] types a restricted linear name if it (a) has no unauthorized roles and (b) has no un*matched* communications (no output or input communication prefixes). Rule [T-snew] types a restricted shared name, without any additional restriction on unauthorized roles. Rules [T-in], [T-out], [T-lsin] and [T-lsout] type input/output actions under the role $\lfloor r\rceil$, with the

authorization set enlarged with (a,r) in case $\lfloor r \rceil$ is $\lfloor r \rfloor$. Notice the typing environment in the conclusion of rules [T-out] and [T-sout] mentions the usage delegated in the communication (via splitting). Rules [T-sin] and [T-sout] state that input and output actions on shared channels are well typed only under authorized roles. The authorization is not performed on shared channels, implying that $\mathtt{pr}_1(\Sigma)$ are linear and not changed under actions on shared channels. Rule [TProc-par] states that the unauthorized pairs in a parallel composition of two processes is the union of unauthorized pairs of the two composed processes. We say that a process P is *well typed* if there are Δ and Γ such that $\Delta;\Gamma \vdash_{\emptyset} P$.

Proposition 1 (Error free) *If P is a well-typed process, then P is not an authorization error.*

We define the reduction relation $\rightarrow$ between behavioral types B and corresponding environments Δ by allowing a synchronized communication prefix to reduce to its continuation $(s \rightarrow r) : l(M).B \rightarrow B$, so as to mimic the respective process behavior, by allowing reduction to occur in a branch of a parallel composition (e.g., $B_1 \rightarrow B_2 \implies B \mid B_1 \rightarrow B \mid B_2$), and by lifting the relation point-wise to environments, embedding reflexivity so as to encompass process reductions involving shared or bound channels (where no reduction in the linear usages of free names is required).

Theorem 1 (Type Preservation) *Let $\Delta;\Gamma \vdash_{\Sigma} P$ for some Δ,Γ,Σ and P. If $P \rightarrow Q$ then there is Δ' such that $\Delta \rightarrow \Delta'$ and $\Delta';\Gamma \vdash_{\Sigma} Q$.*

A direct consequence of type preservation is *protocol fidelity*: every reduction of the process corresponds to a reduction of the types, thus ensuring that the process follows the protocols prescribed by the types. Notice that communication safety (no type mismatches in communications) is entailed by protocol fidelity, which in our case also attests that processes agree in the role when sending and receiving authorizations. Combining freedom from errors and type preservation results we immediately obtain our notion of type safety, which ensures that well-typed processes never reach an error configuration.

Corollary 1 (Type safety) *If P is a well-typed process and $P \rightarrow^* Q$, then Q is not an authorization error.*

4 Related Work and Concluding Remarks

Role-based access control in distributed systems with dynamic access rights was handled in [8] by means of a type system, which ensures security properties. In this calculus, roles assigned to data can be dynamically administered, while role communication between processes was not treated.

Previous works consider security properties, like confidentiality and integrity, in the setting of session calculi. For instance, in [2] session types are extended with correspondence assertions, a form of dependent types which ensures consistency of data during computation. More recently, aspects of secure information flow and access control have been addressed for sessions in [3, 5, 6]. A kind of role-based approach is used in [7], where communication is controlled by a previously acquired reputation.

Similarly to our work, the work [9] consider a typed approach to role-based authorizations, in the setting of service-oriented applications. Differently from our model, assigned roles are initially authorized and communicated data carries information on roles that will use it.

Our contribution is based on the previous work on conversation types [4] and their extension with dynamic assignment of roles to several parties in a concurrent system [1]. We focused on a modular extension of the existing framework, so as to leverage on the previous results, adding the minimal elements so as to identify the specific issues at hand and set the basis for further exploration. We consider the problem of role authorization and authorization passing in an extension of the π-calculus, where communication prefixes are annotated with role authorizations. The underlying calculus allows for the dynamic communication of authorizations. We then extend the conversation type system in which a well-typed

process can never incur in an authorization error. In this way we can statically distinguish processes that are always authorized to communicate on behalf of a role including when authorizations are dynamically passed in messages. As a natural continuation of this study, we aim to extend the present calculus with tools that will enable role de-authorization. For this purpose, we aim to equip the type system with qualified (authorized or unauthorized) roles, instead of unqualified ones. In such a calculus, we could model authorization removal and authorization lending. Moreover, by introducing a partial order into the set of roles, we could control communicated roles with the aim to provide absence of authorization leaks.

Acknowledgments. We are grateful to the anonymous reviewers for their useful remarks. This work was supported by COST Action IC1201: Behavioural Types for Reliable Large-Scale Software Systems (BETTY) via Short-Term Scientific Mission grants (to Pantović and Vieira), and by FCT through project Liveness, PTDC/EIACCO/117513/2010, and LaSIGE Strategic Project, PEstOE/EEI/UI0408/2014 and by grants ON174026 and III44006 of the Ministry of Education and Science, Serbia.

References

[1] Pedro Baltazar, Luís Caires, Vasco T. Vasconcelos & Hugo Torres Vieira (2012): *A Type System for Flexible Role Assignment in Multiparty Communicating Systems*. In Catuscia Palamidessi & Mark Dermot Ryan, editors: *TGC 2012, Lecture Notes in Computer Science* 8191, Springer, pp. 82–96. Available at `http://dx.doi.org/10.1007/978-3-642-41157-1_6`.

[2] Eduardo Bonelli, Adriana B. Compagnoni & Elsa L. Gunter (2005): *Correspondence assertions for process synchronization in concurrent communications*. *J. Funct. Program.* 15(2), pp. 219–247. Available at `http://dx.doi.org/10.1017/S095679680400543X`.

[3] Viviana Bono, Sara Capecchi, Ilaria Castellani & Mariangiola Dezani-Ciancaglini (2011): *A Reputation System for Multirole Sessions*. In Roberto Bruni & Vladimiro Sassone, editors: *TGC, Lecture Notes in Computer Science* 7173, Springer, pp. 1–24. Available at `http://dx.doi.org/10.1007/978-3-642-30065-3_1`.

[4] Luís Caires & Hugo Torres Vieira (2010): *Conversation types*. *Theoretical Computer Science* 411(51-52), pp. 4399–4440. Available at `http://dx.doi.org/10.1016/j.tcs.2010.09.010`.

[5] Sara Capecchi, Ilaria Castellani & Mariangiola Dezani-Ciancaglini (2011): *Information Flow Safety in Multiparty Sessions*. In Bas Luttik & Frank Valencia, editors: *EXPRESS, EPTCS* 64, pp. 16–30. Available at `http://dx.doi.org/10.4204/EPTCS.64.2`.

[6] Sara Capecchi, Ilaria Castellani, Mariangiola Dezani-Ciancaglini & Tamara Rezk (2010): *Session Types for Access and Information Flow Control*. In Paul Gastin & François Laroussinie, editors: *CONCUR, Lecture Notes in Computer Science* 6269, Springer, pp. 237–252. Available at `http://dx.doi.org/10.1007/978-3-642-15375-4_17`.

[7] Pierre-Malo Deniélou & Nobuko Yoshida (2011): *Dynamic multirole session types*. In Thomas Ball & Mooly Sagiv, editors: *POPL*, ACM, pp. 435–446. Available at `http://doi.acm.org/10.1145/1926385.1926435`.

[8] Mariangiola Dezani-Ciancaglini, Silvia Ghilezan, Svetlana Jaksić & Jovanka Pantović (2011): *Types for Role-Based Access Control of Dynamic Web Data*. In Julio Mariño, editor: *WFLP 2010, Lecture Notes in Computer Science* 6559, Springer, pp. 1–29. Available at `http://dx.doi.org/10.1007/978-3-642-20775-4_1`.

[9] Alessandro Lapadula, Rosario Pugliese & Francesco Tiezzi (2007): *Regulating Data Exchange in Service Oriented Applications*. In Farhad Arbab & Marjan Sirjani, editors: *FSEN, Lecture Notes in Computer Science* 4767, Springer, pp. 223–239. Available at `http://dx.doi.org/10.1007/978-3-540-75698-9_15`.

[10] Davide Sangiorgi & David Walker (2001): *The Pi-Calculus - a theory of mobile processes*. Cambridge University Press.